

Article (refereed)

Keef, Caroline; **Svensson, Cecilia**; Tawn, Jonathan A.. 2009
Spatial dependence in extreme river flows and precipitation for
Great Britain. *Journal of Hydrology*, 378 (3-4). 240-252.
[10.1016/j.jhydrol.2009.09.026](https://doi.org/10.1016/j.jhydrol.2009.09.026)

Copyright © 2009 Elsevier B.V. All rights reserved

This version available <http://nora.nerc.ac.uk/8814/>

NERC has developed NORA to enable users to access research outputs wholly or partially funded by NERC. Copyright and other rights for material on this site are retained by the authors and/or other rights owners. Users should read the terms and conditions of use of this material at <http://nora.nerc.ac.uk/policies.html#access>

This document is the author's final manuscript version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this and the publisher's version remain. You are advised to consult the publisher's version if you wish to cite from this article.

www.elsevier.com

Contact CEH NORA team at
noraceh@ceh.ac.uk

Spatial dependence in extreme river flows and precipitation for Great Britain

Caroline Keef*, Cecilia Svensson[†] and Jonathan A. Tawn[‡]

September 1, 2009

Abstract

For the co-ordination of flood mitigation activities and for the insurance and re-insurance industries, knowledge of the spatial characteristics of fluvial flooding is important. Past research into the spatio-temporal risk of fluvial flooding is restricted to empirical estimates of risk measures and hence estimates cannot be obtained for return periods longer than the length of the concurrent data at the sites of interest in the sample. We adopt a model-based approach which describes the multisite joint distribution of daily mean river flows and daily precipitation totals. A measure of spatial dependence is mapped across Great Britain for each variable separately. Given that an extreme event has occurred at one site, the measure characterises the extent to which neighbouring locations are affected. For both river flow and precipitation we are able to quantify how events become more localised in space as the return periods of these events get longer at a site of interest. For precipitation, spatial dependence is weaker in the upland areas of Great Britain. For river flows the major factor affecting spatial dependence appears to be differences in catchment characteristics with areas with diverse catchments exhibiting lower levels of dependence.

Keywords: extreme value theory, multivariate extreme values, precipitation, river flow, spatial dependence.

* corresponding author: JBA Consulting, South Barn, Broughton Hall, Skipton, North Yorkshire, BD23 3AE, UK. Email: caroline.keef@jbaconsulting.co.uk

[†] CEH Wallingford, Wallingford, Oxon, OX10 8BB, UK. Email: csve@ceh.ac.uk

[‡] Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4YF, UK. Email: j.tawn@lancaster.ac.uk

1 Introduction

For the co-ordination of flood mitigation activities and for the insurance and re-insurance industries, knowledge of the spatial characteristics of fluvial flooding is important. This study focuses on the spatial dependence in extreme river flows with the aim being to map a measure of spatial dependence of extreme values over Great Britain. We analyse data series of daily mean flows as these exist in long concurrent records of high quality. To provide an insight into how differences in regional meteorology and catchment characteristics influence spatial flood risk we also analyse the spatial dependence in extreme daily precipitation totals, and carry out an analysis of the dependence between catchment characteristics and a measure of spatial dependence in extreme river flows.

The spatial risk of flooding (or extreme precipitation) at a set of m gauged sites, denoted $\Delta = \{1, \dots, m\}$, can be described in many ways. For a spatial risk measure to have generic applicability it is important that it is not specific to the set of gauged sites. Therefore we measure how extreme an event is relative to the distribution at each site, i.e. on the probability scale. The distributions of the variables will be different over all m sites but the probability scale is common over all m sites. Thus levels at two sites are viewed to be equally extreme if they both correspond to the p th quantile of the variable at each site even though their actual flow levels or precipitation amounts can differ considerably. It is helpful, however, for interpretation to present results in terms of return periods of events instead of quantiles of the variable. For a process which is observed daily, with extreme events lasting on average k days, the p^{th} quantile corresponds to the T year event, with

$$T = \frac{k}{365(1 - p)}, \quad (1.1)$$

provided $T > 1$ year to avoid issues of seasonality. As the return period depends on k , which will be different for different processes, we typically focus on the probability scale p but for applications we will present results in terms of return periods.

Our approach is to derive an estimate of spatial dependence centred at each of a network of gauged sites indexed by Δ . For each site $i \in \Delta$, let $\Delta_i(d) \subseteq \Delta$ denote the sites in Δ , excluding site i , that are within d km of site i . We will evaluate the spatial dependence between the variables

1 at the sites in $\Delta_i(d)$ conditional upon an extreme event, i.e. an event which exceeded the T
2 year return level, occurring at site i . Our main measure of this conditional spatial dependence
3 will be $N_i(p, d)$, the expected proportion of sites in $\Delta_i(d)$ that exceed their p th quantile (or
4 equivalently T year return level) during an event in which site i exceeds its p th quantile (T year
5 return level), we define $N_i(p, d)$ precisely in Section 3. In this paper we estimate the measure
6 $N_i(p, d)$, and related dependence summaries, for a range of T (equivalently p), including very
7 large T , and for a range of d .

8 This raises a number of issues including: How are suitable events in spatial series identified?
9 How is $N_i(p, d)$ estimated, particularly when T (equivalently p) is large as empirical estimates
10 fail when T exceeds the observed amount of overlapping data series? If a model for dependence
11 is to be used can it be selected to be sufficiently general to apply to all possible combinations of
12 sites? How are missing data to be handled as there are limited concurrent data when many sites
13 are considered simultaneously? And how are confidence intervals for the estimate evaluated?

14 This is essentially a multivariate extreme values problem: given that an extreme event
15 occurred at site i we are concerned with the occurrence of extreme values at sites in $\Delta_i(d)$ during
16 the same event. Though a number of measures of bivariate and multivariate dependence have
17 been developed in the hydrological literature to describe dependence in extreme values, see Dales
18 and Reed (1989) and Svensson and Jones (2002), they are empirical. Hence, they are restricted
19 to measuring dependence at return periods at most as long as the length of the data series. They
20 are also dependence measures for specific subclasses of (multivariate extreme value) distributions
21 so will produce biased estimates when these distributional choices are inappropriate. Another
22 feature of these methods is that for each extreme event, observations are required at all sites,
23 so events with missing components at any site are removed.

24 We use the methodology developed by Keef et al. (2009) which extends the approach devel-
25 oped by Heffernan and Tawn (2004). This method adopts an asymptotically justified dependence
26 model over sites and time which accounts for time lags between variables at different sites. The
27 approach has a number of benefits: it allows extrapolation to events of any return period; it
28 compares favourably with other available methods in terms of its ability to model a broad range
29 of dependence structures; it is applicable to a large number of sites; and it accounts for missing
30 data. In using these methods we make the assumption that fluvial flooding in the winter (higher

flow) months has the same spatial characteristics as fluvial flooding in the summer (lower flow) months. The simplest alternative to this assumption is to split the data up into higher and lower flow months. The difficulty with this alternative is that there are very few floods in lower flow months. In Environment Agency (2009) almost no difference was found between estimating $N_i(p, d)$ using data from the full year and only data from the higher flow months.

The paper is structured as follows: Section 2 presents the network of data used in the study together with some information about the key covariates; in Section 3 a summary is given of the methods of Keef et al. (2009); in Section 4 maps over Great Britain of spatial dependence in daily river flows and precipitation are presented and interpreted; and in Section 5 conclusions are summarised.

2 Data selection and key covariates

For our analysis of daily precipitation data we use data from 256 raingauges across Great Britain (GB) provided by the UK Met Office. For the daily mean river flow we use data from a set of 271 stations in GB provided by the National River Flow Archive (NRFA) (Marsh and Hannaford, 2008). For the river flow series the selected stations have on average 40 years of data with a minimum of 20 years and maximum of 60 years. For the raingauges almost all of the selected sites have 40 years of data, mostly corresponding to the same period, since 1961, so have long spans of overlapping series from gauge to gauge.

Though many more stations/gauges exist we selected sites which are recognised as having high quality data, long records with limited missing data, that provide a good spatial coverage of GB and that have long periods of overlap with each other. The stations were also selected to have as homogeneous records as possible; this selection avoided stations influenced by the construction of new dams and flood defence systems or stations with catchments which have experienced major changes in land use. When stations that are very close to each other (less than around 25 km apart) were identified as appropriate only the station with the longest records was selected.

Figures 1 and 2 shows the locations of the selected river flow stations and raingauges respectively. In some areas of GB it was not possible to obtain a network of similar density of good quality river flow records. For example, in north-west Scotland and in the East Anglian Fens

1 the density of the NRFA gauging stations is more sparse relative to the rest of GB. Due to the
2 large number of raingauges to choose from we were able to achieve a better spatial coverage
3 of the selected raingauges. However, the density is still not uniform over GB, with the density
4 being notably lower in the Scottish Highlands and Wales.

5 Figure 1 shows the baseflow index (BFI) and areas of the selected catchments. The BFI
6 may be thought of as a measure of the proportion of the total discharge that derives from
7 stored sources, and is a way of indexing the catchment geology. Permeable rocks and soils can
8 store more water, resulting in a higher baseflow and more sustained flow during periods of dry
9 weather. Impervious clay catchments with minimal lake and reservoir storage typically have
10 BFIs in the range 0.15 – 0.35, whereas a chalk stream may have a BFI greater than 0.9 (Marsh
11 and Hannaford, 2008). Figure 2 shows the raingauge altitudes and standard annual average
12 rainfalls (SAAR). The mapped SAAR values reveal a clear rainshadow effect. Precipitation in
13 northern and western Britain is predominantly of frontal origin, tending to fall in south-westerly
14 airflow, e.g. Manley (1970). Orographic enhancement of precipitation over mountainous areas
15 therefore mean that south- and west-facing slopes receive considerably more precipitation than
16 areas to the east of the hills. In contrast, widespread and persistent precipitation in eastern
17 GB tends to occur in the onshore winds on the north side of depressions passing eastwards on
18 a more southerly track, e.g. Mayes and Sutton (1997).

19 Two sets of four river flow gauges each were used as an initial test dataset when developing
20 the methods, and are used to demonstrate particular features of the dependence in Section 4.
21 The gauges are located in south-east Scotland and in the Thames catchment in south England.
22 They were chosen to contain only a small amount of missing data, and to represent different
23 catchment characteristics. Table 1 shows the catchment area and BFI for the catchments of these
24 gauging stations, it shows that the Scottish rivers generally have a lower BFI than the Thames
25 rivers. They also tend to be hillier and smaller in catchment size. Nearly all the catchments are
26 essentially rural. The exception is the Thames at Kingston, the most downstream gauge on the
27 Thames, which has mixed land use including urban areas.

3 Methods

3.1 Strategy

Consider a set Δ of gauging sites, with hydrological variable X_i at gauging site i . The joint distribution of any multivariate random variable, $\mathbf{X} = \{X_i, i \in \Delta\}$, where each $X_i, i \in \Delta$ has a continuous distribution function, can be separated into marginal distribution functions, $F_i, i \in \Delta$ where $F_i(x) = \Pr(X_i \leq x)$, and the dependence structure. This separation of the joint distribution can be expressed as follows

$$\Pr(X_1 \leq x_1, \dots, X_m \leq x_m) = C\{F_1(x_1), \dots, F_m(x_m)\} \quad (3.1)$$

where C is a unique function, known as the copula, which determines the dependence structure of \mathbf{X} , see Joe (1997) and Nelsen (1999). The copula formulation separates the joint distribution into the m marginal distribution functions and a joint distribution function C for the variables on a common marginal distribution. In expression (3.1), C is the joint distribution function for uniform $[0, 1]$ variables, however the choice of common marginal distribution does not matter and so in different studies different common marginal distribution forms are selected for the convenience of the problem.

For the construction of our spatial dependence measure we are interested in the dependence structure of \mathbf{X} only, and so we focus on the copula component of expression (3.1). Therefore we initially transform our variable \mathbf{X} , e.g. river flow or precipitation, to a common marginal distribution. For the purposes of modelling extreme values it is best to use standard Gumbel margins as this scale induces the most linearity in the dependence (Heffernan and Tawn, 2004). Therefore we transform \mathbf{X} to \mathbf{Y} , so that \mathbf{Y} has standard Gumbel marginal distributions, through the transformation $Y_i = -\log\{-\log[F_i(X_i)]\}$ so that $\Pr(Y_i \leq y) = \exp[-\exp(-y)]$, for $-\infty < y < \infty$ and $i \in \Delta$. We denote the p th quantile of Y_i , for each i , by y_p . As all $F_i, i \in \Delta$, are unknown we use the empirical distribution function of X_i to estimate them, so it is only the ranks of \mathbf{X} that we use and not the values themselves. An important consequence of this approach is that the effect of measurement error (including random and systematic errors) on flows is small for this analysis.

The aim is to be able to estimate features of the distribution of a multivariate random

variable Y when Y_i is large under the assumption that observations on Y are from a stationary time series. Let Y_{-i} denote Y with Y_i removed. The strategy is first to model the marginal distribution of Y_i , which is given due to our choice of Gumbel distribution. Then we model the conditional distribution of $Y_{-i}|Y_i = y$ for large values of y . Here the conditional distribution is the distribution of all the other elements of Y given that Y_i is fixed equal to y . We propose a model motivated by asymptotic probabilistic theory for this conditional distribution as $y \rightarrow \infty$. We assume that this model will be appropriate for all values of y above a high threshold u (in this paper we take $u = y_{0.975}$ throughout), and use the observations of Y with $Y_i > u$ to fit this model. Having fitted the conditional distribution the final stage is the estimation of features of the fitted distribution of Y when $Y_i > v$ where $v \geq u$.

The key step in all this process is the modelling of the conditional distribution of $Y_{-i}|Y_i = y$ for large y as all subsequent aspects of the inference hinge on this step. This section focuses on this aspect. Heffernan and Tawn (2004) provide asymptotic theory which provides the basis for a model for the behaviour of the conditional distribution of $Y_{-i}|Y_i$, when Y_i is large, and they perform inference for this model under the assumption that observations of Y are independent, identically distributed, and without missing values. Here we will focus on a special case of this model which appears appropriate for our applications to fluvial and precipitation extreme value modelling. Specifically, we assume that all variables are non-negatively dependent. Full details of the theory on which the model is based, the associated conditional models, and a discussion of when it is appropriate to make simplifying assumptions such as those we make here are given in Heffernan and Tawn (2004) and Keef et al. (2009). In Sections 3.2 and 3.3 we describe the associated conditional model for extreme values in bivariate and multivariate cases respectively. Then in Sections 3.4 and 3.5 we outline the extensions proposed by Keef et al. (2009) to cover temporal dependence and missing data. Our spatial dependence measure is presented in Section 3.6 and the methods used to evaluate uncertainty are given in Section 3.7.

3.2 Bivariate model

Consider a high threshold u and any pair of sites $i, j \in \Delta$. When the variables are non-negatively dependent, the simplified version of Heffernan and Tawn's conditional model for $Y_j|Y_i = y$, when

1 $y > u$, is

$$Y_j = \alpha_{j|i}y + y^{\beta_{j|i}}Z_{j|i} \quad (3.2)$$

2 where $0 \leq \alpha_{j|i} \leq 1$ and $-\infty < \beta_{j|i} < 1$ are parameters and $Z_{j|i}$ is a random variable with
 3 non-zero mean, independent of Y_i , and with distribution function $G_{j|i}(z) = \Pr(Z_{j|i} \leq z)$. The
 4 restrictions of this model are discussed in Heffernan and Tawn (2004). The formulation of the
 5 model through expression (3.2) is a regression model of Y_j upon Y_i , for Y_i being above the
 6 threshold u , with conditional mean and variance given by

$$E(Y_j|Y_i = y) = \alpha_{j|i}y + y^{\beta_{j|i}}E(Z_{j|i}) \text{ and } \text{Var}(Y_j|Y_i = y) = y^{2\beta_{j|i}}\text{Var}(Z_{j|i})$$

7 for $y > u$; with the quantities $Z_{j|i}$ being equivalent to the standardised residuals in a classical
 8 regression model.

9 The parameters $\alpha_{j|i}$ and $\beta_{j|i}$ describe the strength of dependence between (Y_i, Y_j) when Y_i
 10 is large. Unlike classical regression, $E(Z_{j|i}) = 0$, so the interpretation of $\alpha_{j|i}$ and $\beta_{j|i}$ is not
 11 as straight-forward as in a standard linear regression model. The parameter $\alpha_{j|i}$ describes the
 12 overall strength of dependence between the two variables, as $\alpha_{j|i}$ increases the overall strength of
 13 dependence between Y_i and Y_j increases. In contrast $\beta_{j|i}$ describes how the dependence changes
 14 with Y_i ; for positive values of $\beta_{j|i}$ the variance of $Y_j|Y_i = y$ increases as y increases.

15 To get a clearer impression of the roles of $\alpha_{j|i}$ and $\beta_{j|i}$ consider the two examples given in
 16 Figure 3 showing samples of (Y_i, Y_j) when $Y_i > 4.5$, corresponding to when Y_i is in the top 1%
 17 of the standard Gumbel distribution. In the left hand plot both large and small values of Y_j can
 18 occur with large Y_i , with large values of Y_j occurring with large values of Y_i more often than if
 19 (Y_i, Y_j) were independent. In the right hand plot, only large values of Y_j occur with large values
 20 of Y_i . This illustrates that large values of Y_j can occur with large values of Y_i if either $\alpha_{j|i}$ is
 21 large or if $\beta_{j|i}$ is large. However, small and large values of Y_j occur with large Y_i only when $\beta_{j|i}$
 22 is large.

23 The two samples shown in Figure 3 illustrate different degrees of extremal dependence. To
 24 help clarify the differences the associated probabilities of

$$\Pr(Y_j > v|Y_i > v) \quad (3.3)$$

are also given in Figure 3 for $v > 4.5$. In each case we compare with the true value of probability (3.3), which can be evaluated as we know the joint distribution of Y that was used to simulate these data. We see that both the true probability and its empirical estimate is larger for the second sample than for the first indicating that the second parameter set $(\alpha_{j|i}, \beta_{j|i}) = (0.8, 0.1)$, leads to a stronger form of positive extremal dependence.

The strongest form of extremal dependence is asymptotic dependence which occurs when $\alpha_{j|i} = 1$ and $\beta_{j|i} = 0$ as probability (3.3) tends to a non-zero constant as $v \rightarrow \infty$. In contrast, when $\alpha_{j|i} < 1$ then probability (3.3) tends to zero as $v \rightarrow \infty$, with this case being termed asymptotic independence. Asymptotic dependence and asymptotic independence are discussed in Coles et al. (1999) and their relation to the Heffernan and Tawn model is discussed in Heffernan and Tawn (2004) and Keef et al. (2009). Under independence of Y_j and Y_i then $P(Y_j > v | Y_i > v) = P(Y_j > v) \approx e^{-v}$, for large v .

Estimates $\hat{\alpha}_{j|i}$ and $\hat{\beta}_{j|i}$ of the parameters $\alpha_{j|i}$ and $\beta_{j|i}$ are obtained using likelihood methods under a working assumption that $Z_{j|i}$ follows a Normal distribution. Specifically we assume that each random variable $Y_j | Y_i = y$, for $y > u$, $j \in \Delta \setminus \{i\}$ has mean $\mu_{j|i}(y)$, and standard deviation, $\sigma_{j|i}(y)$, respectively given by

$$\mu_{j|i}(y) = \alpha_{j|i}y + \mu_{j|i}y^{\beta_{j|i}}$$

$$\sigma_{j|i}(y) = \sigma_{j|i}y^{\beta_{j|i}}.$$

This enables us to estimate the parameters $\alpha_{j|i}$, $\beta_{j|i}$, $\mu_{j|i}$ and $\sigma_{j|i}$ by applying a numerical maximisation procedure to the following function

$$-\sum_{t=1}^{t_j} \left[\log\{\sigma_{j|i}(y_{i,t})\} + \frac{1}{2} \frac{(y_{j,t} - \mu_{j|i}(y_{i,t}))^2}{\sigma_{j|i}(y_{i,t})} \right],$$

where t_j is the number of observations of y_j when $y_i > u$. This method of parameter estimation does not introduce a noticeable bias in the parameter estimates. Heffernan and Tawn (2004) found some restrictions on $G_{j|i}$ but that these were insufficient to justify using any parametric family, so the distribution $G_{j|i}$ is estimated using the empirical distribution of the estimated standardised residuals

$$\hat{z}_{j|i} = \frac{Y_j - \hat{\alpha}_{j|i}Y_i}{Y_i^{\hat{\beta}_{j|i}}} \text{ for } Y_i > u. \quad (3.4)$$

The key to the choice of the threshold u is that for $Y_i > u$ then Y_i and $Z_{j|i}$ are independent. The theory suggests that there should always be some level u above which independence is an appropriate assumption. Therefore u needs to be selected large enough to achieve this independence, with a test of independence guiding this choice, but u needs to be small enough to ensure there are sufficient data with $Y_i > u$ to be able to estimate the model reliably.

To derive the distribution of any features of the joint distribution of (Y_i, Y_j) , such as maximum of (Y_i, Y_j) , or $Y_i + Y_j$ we need to study the behaviour of function h of (Y_i, Y_j) . The distribution of any functional, $h(Y_i, Y_j)$ of (Y_i, Y_j) with $Y_i > v$, for $v > u$ can be derived by Monte Carlo methods by simulating samples from the fitted distribution $(Y_i, Y_j)|Y_i > v$ and evaluating h for each replicate. To do this first simulate $Y_i > v$ giving a value y^* say and second simulate $Y_j|Y_i = y^*$. As a consequence of the Gumbel marginal distribution for Y_i , the simulation from $Y_i|Y_i > u$, when u is large, corresponds to simulating from a standard exponential distribution shifted by u . The key to simulating from the conditional $Y_j|Y_i = y^*$ is that under the conditional model (3.2) Y_i and $Z_{j|i}$ are independent. Hence independently of Y_i , the variable $Z_{j|i}$ is simulated from the estimated distribution $G_{j|i}$, i.e. $Z_{j|i}$ is sampled from replicates of $\hat{Z}_{j|i}$, and Y_j is derived using expression (3.2). Using a similar strategy the distribution of the functional can be obtained under independence and perfect dependence by repeatedly simulating Y_j independently of Y_i or taking $Y_j = Y_i$. These cases are sometimes helpful to assess the impact of dependence. The sample size of the simulated sample is taken to be sufficiently large that the Monte Carlo component of the uncertainty of the empirical estimate is very small. The level v , $v > u$, can be arbitrarily large and so the proposed strategy provides estimates of features about the distribution of Y within the observed tail of Y_i through to extrapolation beyond the maximum Y_i observation.

Taking $u = 4.5$ we illustrate using this method for the simulated data to estimate $\Pr(Y_j > v|Y_i > v)$, for $v > u$ in Figure 3. Using the simulation method for evaluating an estimate of $\Pr(Y_j > v|Y_i > v)$ under the fitted model as described in Section 3.1, we see in Figure 3 that the model-based estimate matches well with both the empirical and true values in the range of the data but continues to give a smooth extrapolation to the empirical estimate and as we know here the true value we can also see it gives a good estimate of the true value of $\Pr(Y_j > v|Y_i > v)$ for v beyond the largest observed Y_i value. This may not be too surprising as the basic model

1 formulation is correct, but recall that the α and β parameters needed to be estimated and only
 2 a nonparametric estimate of the form of the distribution of $Z_{j|i}$ was used.

3 3.3 Multivariate model

4 Now consider the extension of equation (3.2) to the multivariate case. The simplified Heffernan
 5 and Tawn model is

$$Y_{-i} = \alpha_{|i}y + y^{\beta_{|i}}Z_{|i} \quad (3.5)$$

6 where $0 \leq \alpha_{|i} \leq 1$ and $-\infty < \beta_{|i} < 1$ are parameter vectors of dimension $|\Delta| - 1$ and $Z_{|i}$ is a
 7 $|\Delta| - 1$ dimensional random variable with joint distribution function $G_{|i}$ and is independent of
 8 Y_i . Here vector algebra is to be interpreted as componentwise.

9 Model (3.5) is a multivariate regression model with for every $j \in \Delta$ with $j \neq i$ the j th
 10 marginal component being identical to the bivariate model (3.2). By parametrically modelling
 11 the dependence of each component of Y_{-i} on Y_i we remove some of the dependence between the
 12 Y_{-i} variables. An explanation of this is that if all variables in Y_{-i} are positively associated with
 13 Y_i , they will all tend to be large or small together. This tendency is mostly captured by the α
 14 and β parameters. Once the common dependence on Y_i is accounted for, the residual variation,
 15 determined by the joint distribution $G_{|i}$ of the standardised multivariate residuals $Z_{|i}$, is then
 16 modelled. We anticipate that $Z_{|i}$ will have weaker dependence than Y_{-i} ; for example if Y_{-i}
 17 are conditionally independent given Y_i then the Y_{-i} can possess arbitrarily strong dependence
 18 but $Z_{|i}$ will be independent.

19 This residual dependence is captured by the inclusion of a non-parametric model for $G_{|i}$
 20 with the empirical joint distribution being used. This corresponds to empirically estimating
 21 the standardised residuals given by expression (3.4) for all $j \in \Delta$, to provide an estimate $\hat{Z}_{|i}$
 22 for simultaneous observations. Thus our model for $G_{|i}$ has a dependence structure between
 23 the components of $Z_{|i}$ which is identical to the sample dependence structure in their observed
 24 values. The assumption of a non-parametric model for $G_{|i}$ is not as restrictive for extrapolation
 25 of the distribution of Y for large Y_i as in some other modelling contexts. The reason for this
 26 is that the key elements to extrapolation of this distribution are the marginal distribution of Y_i
 27 and the expectation and variance structure of the conditional distribution in the regression, all
 28 of which are modelled parametrically.

3.4 Multivariate-temporal model

Let $Y_{j,t}$ denote variable Y_j on day t . By assuming stationarity of the multivariate time series we are able to model the distribution of the variable at site j at a lag τ in relation to an extreme value at site i , i.e. the conditional distribution $Y_{j,t+\tau}|Y_{i,t}$, for $Y_{i,t} > u$, for any t, τ and j using the methods of Heffernan and Tawn (2004). The methods extend directly with additional notation to denote the lag, i.e. parameters $\alpha_{j|i}^{(\tau)}$ and $\beta_{j|i}^{(\tau)}$ and associated variable $Z_{j|i}^{(\tau)}$. An exploratory summary measure of bivariate dependence that we use in this case is

$$P_{j|i}^{(\tau)}(p) = \Pr(Y_{j,t+\tau} > y_p | Y_{i,t} > y_p), \quad (3.6)$$

so, $P_{j|i}^{(\tau)}(p)$ is the probability that Y_j exceeds y_p at lag τ from Y_i exceeding y_p . As $p \rightarrow 1$ then $P_{j|i}^{(\tau)}(p) \rightarrow 0$ if $\alpha_{j|i}^{(\tau)} < 1$ but tends to a non-zero constant if $\alpha_{j|i}^{(\tau)} = 1$. We consider two special features of this measure, τ_{\max} , the value of τ for which $P_{j|i}^{(\tau)}(p)$ is a maximum, and the value of the conditional probability at this lag, $P_{j|i}^{(\tau_{\max})}(p)$.

Similarly the methods can be applied to lagged values of a single series, i.e. $Y_{i,t+\tau}|Y_{i,t}$, giving a useful measure for clustering of extremes in a single series to be $P_{i|i}^{(\tau)}(p)$, and a multivariate extension of measure (3.6) is studied by Keef et al. (2009).

Because the peak values at the different sites may occur some time apart a dependence measure of more practical value is given by

$$Q_{j|i}^{(L)}(p) = \Pr \max_{\tau \in A_L} Y_{j,t+\tau} > y_p | Y_{i,t} > y_p \quad (3.7)$$

where $A_L = \{-L, -L+1, \dots, 0, \dots, L\}$, so $Q_{j|i}^{(L)}(p)$ is the probability that Y_j has at least one exceedance of y_p within L days of Y_i exceeding y_p . Here the choice of the maximal lag L depends on the temporal dependence at each site and the cross-series lags which we identify in Sections 4.2 and 4.3. As $p \rightarrow 1$ then $Q_{j|i}^{(L)}(p) \rightarrow 0$ if $\alpha_{j|i}^{(\tau)} < 1$ for all $\tau \in A_L$ but tends to a non-zero constant if there is at least one lag $\tau \in A_L$ for which $\alpha_{j|i}^{(\tau)} = 1$.

3.5 Handling missing data

Keef et al. (2009) show full details of the method for an arbitrary missing data pattern. To illustrate the approach for handling missing data here consider three sites ($m = 3$), with site 1

being the conditioning site, all values of Y being observed at all times except time t for which $(Y_{1,t}, Y_{2,t})$ are observed, with $Y_{1,t} > u$, but $Y_{3,t}$ is missing. It follows that $Z_{3|1}$ at time t , denoted $Z_{3|1,t}$, is missing but $Z_{2|1,t}$ is observed. The strategy is to replace the missing value of $Z_{3|1,t}$ by a sample from the conditional distribution of $Z_{3|1,t}|Z_{2|1,t}$. For the general case of m sites the conditional distribution of interest is the distribution of all the missing elements of $Z_{|i}$ given all the observed elements of this vector.

The form of the conditional distribution arises from making an assumption about the copula family for the joint distribution of $Z_{|i} = (Z_{2|1}, Z_{3|1})$. Specifically, a Gaussian copula is assumed (Joe, 1997) with its parameters estimated from the observed pairs of $Z_{|i}$. This assumption is only made in estimating the conditional distribution of the missing values. The method of estimating the distribution of the observed values does not change. Consequently, estimation of features of the joint distribution can be derived using the simulation methods described in Section 3.1 with missing components simply replaced by sample values from the conditional distribution of the missing component given the observed component.

3.6 Spatial dependence measures

We can now specify the spatial dependence measure $N_i(p, d)$ that we introduced in Section 1. Given that $Y_{i,t} > y_p$, then $N_i(p, d)$ is the expected proportion of stations in $\Delta_i(d)$ that exceed y_p for at least one lag in the same event (i.e. within a lag of L days). Mathematically this is expressed as

$$N_i(p, d) = \frac{E(\#\{\mathbf{j} \in \Delta_i(d) : \max_{\tau \in A_L} Y_{\mathbf{j}, t+\tau} > y_p\} | Y_{i,t} > y_p)}{\#\{\mathbf{j} \in \Delta_i(d)\}}, \quad (3.8)$$

where $0 \leq N_i(p, d) \leq 1$ with larger $N_i(p, d)$ corresponding to stronger extremal dependence in the region around site i . By estimating the expected proportion of sites instead of the expected number of sites we aim to overcome the property that sites are not evenly distributed so that the number of neighbouring sites varies from site-to-site.

Changes in $N_i(p, d)$ with p increasing is an indication of how the spatial dependence varies as events become more extreme at site i . If $\alpha_{j|i}^{(\tau)} < 1$ for all $\tau \in A_L$ and all $\mathbf{j} \in \Delta_i(d)$ then $N_i(p, d) \rightarrow 0$ as $p \rightarrow 1$. So it is possible to have strong spatial dependence between sites for moderate return periods but for larger return periods the dependence weakens or becomes so

weak that all the very largest extreme events can only occur at single sites in the network. We will also assess how $N_i(p, d)$ changes with d , we expect it to decrease as d is increased as the additional gauging stations that are included will tend to be more widely separated and so are likely to have less dependence between them.

3.7 Uncertainty

As our model is semi-parametric with a large number of parameters, $2 \sum_{i=1}^m |\Delta_i(d)|$, a natural approach to evaluating uncertainty is to use the block bootstrap. We explored a number of ways for selecting blocks to correspond to independent and identical events (Keef, 2007) but concluded that for a generic application of the methods one-year blocks were ideal as they accounted for both the longer range dependence experienced at some sites and seasonality at all sites. These one-year blocks were chosen to run from 1st August to 31st July to minimise the chance of splitting an extreme fluvial event into different blocks.

4 Results and discussion

4.1 Introduction

In this section we show the results of estimating the dependence measures of Section 3.4, namely $P_{j|i}^{(r)}(p)$, τ_{\max} , $P_{j|i}^{(\tau_{\max})}(p)$ and $Q_{j|i}^{(L)}(p)$, for flow stations and precipitation gauges in GB. Throughout we take $y_p = u$ so $p = 0.975$. As discussed in Section 3.1 we estimate all these features, and $N_i(p, d)$, of the joint distribution by using empirical estimates derived from a large sample of simulated values generated from the fitted distribution. The spatial dependence analysis presented focuses on areas with radii, d , of 30 and 60 km. See Keef (2007) for more details on results for d of 90 and 120 km.

In all cases we only present results for the model-based estimator developed in Section 3. However, as in Figure 3, we assessed the fit of our models by comparing the model-based estimates with empirical estimates over the range for which the empirical estimates are reliable and found a good agreement. Examples of such comparisons and more general diagnostic procedures for these models are given by Heffernan and Tawn (2004) and Keef et al. (2009).

4.2 Temporal dependence in individual series

To illustrate the use of the dependence measures introduced in Section 3.4 we first apply the methods to explore the temporal dependence in river flow extremes at a single site. The extremal dependence in river flows was estimated using $P_{i|i}^{(\tau)}(p)$, equation (3.6). A value of $P_{i|i}^{(\tau)}(p) = 0.5$ means that there is a 50% chance that, conditional on variate Y_i being large (above the p^{th} quantile) on any given day t , then it will also be large on day $t + \tau$. Figure 4 shows the results for $|\tau| \leq 50$ days, for the Tweed at Peebles and the Thames at Eynsham. The faster responding Tweed catchment shows a relatively quick drop in dependence over a few days, whereas in the larger, flatter, more permeable Thames catchment the dependence recedes more gradually over a period of just over 10 days.

4.3 Bivariate temporal dependence

When estimating the probability of several locations experiencing extreme events at the same time, there is a need to define what simultaneous events are. For example, river flows are unlikely to peak at exactly the same time in response to the same rainfall event, because of the different catchment characteristics. Therefore, the lagged extremal dependence measure $P_{j|i}^{(\tau)}(p)$, given by equation (3.6), was estimated for different pairs of gauging stations indexed by i and j for a range of lags, τ .

Figure 5 shows two examples of bivariate temporal dependence in river flows, between the Tweed at Kingledores and the Earn at Kinkell Bridge, and between the Thames at Kingston and the Lambourn at Shaw. For the first pair, the maximum extremal dependence occurs for a lag of 0 days between the series. However, because of the very slow response of the Lambourn, there is no well-defined maximum for the Thames-Lambourn pair. Instead the dependence increases slowly up to a lag of about 18 days, after which it levels off. It is likely that the measured dependence between the Thames-Lambourn pair reflects the shared seasonality between this pair. In particular the flow regime for the Lambourn at Shaw is dominated by seasonal behaviour. This suggests that for some very slowly responding catchments the assumption of stationarity may not be an acceptable assumption even taken for sub periods of a year. For our selected set of gauges, 96% of the river flow gauge-pairs have estimates of extremal dependence that peak within a lag of 3 days, i.e. $|\tau_{\max}| \leq 3$. The station pairs that do not have peak dependence

within 3 days are those for which at least one of the pair is a very slowly responding catchment. Thus we considered river flow events to be simultaneous only if they occur within 3 days of an extreme observation at the conditioning gauge. Hence, we take $L = 3$ in the evaluation of $N_i(p, d)$ for river flows. For precipitation we estimated $N_i(p, d)$ for both $L = 0$ and $L = 1$. We found no noticeable difference for these maximum lags so we take $L = 0$ for the whole analyses.

4.4 Spatial dependence in precipitation

Figure 6 shows estimates of the spatial dependence measure $N_i(p, 30)$ for a range of p , and Figure 7 shows $N_i(p, 60)$ for one of these values of p , with its associated 95% confidence interval. The median numbers of neighbours considered were 2 for $d = 30$ km and 11 for $d = 60$ km. The return periods were calculated by taking $k = 1$ in expression (1.1), we estimated this measure for p equal to 0.975, 0.995, 0.999, 0.9995, 0.9999, 0.99995, 0.99999 and 0.999995 which correspond to the $T = 0.1, 0.5, 3, 5, 30, 55, 274$ and 548 year return periods and for $d = 30$ and 60 .

Both Figures 6 and 7(b) show that spatial dependence is stronger in south-east GB than in the mountainous north and west. This broadly agrees with the results of Dales and Reed (1989), who investigated spatial dependence in annual maximum daily precipitation. The lower dependence in the hilly areas may reflect that orographic precipitation varies spatially at the same scale as the topography, and particularly across major topographical barriers. Svensson and Jones (2002) investigated extremal dependence between pairs of rain-gauges. Using continuous series of observed daily precipitation, they found that dependence breaks down across topographical barriers. The lack of topographic variation in the south and east of the country is likely to be the reason that dependence in heavy precipitation is stronger over a larger area.

Figure 6 shows that $N_i(p, 30)$ decreases with increasing return period T for all T , corresponding to increasing p . This suggests that extreme precipitation is asymptotically independent over sites within $d = 30$ km in the rain-gauge network. For $T = 0.1$ years Figure 6(a) shows $N_i(p, d)$, is between 0.4 and 0.6 for the bulk of GB. For $T = 55$ years, shown in Figure 6(c), the dependence drops to about 0.2–0.4. The decrease in spatial dependence with increasing return period seems reasonable because of the localised character of heavy events. Weak evidence for such a decrease in dependence was also found by Dales and Reed (1989), using only annual maxima data. However, Dales and Reed (1989) were not able to quantify this changing dependence or

extrapolate it within their framework of annual maxima analysis.

Comparing $N_i(p, 30)$ and $N_i(p, 60)$ for p corresponding to $T = 55$ years in Figures 6(c) and 7(b), it can be seen that spatial dependence decreases with increasing d as expected. The strongest dependence occurs where $N_i(p, 30)$ exceeds 0.4 at several locations in the south and east. This means that if a gauge experiences heavy precipitation, more than 40% of gauge locations within 30 km also receive heavy precipitation. In the north and west $N_i(p, 30)$ drops to around 0.2. The corresponding values for $N_i(p, 60)$ are 0.3 for the southeast and 0.1 for the north-west. The spatial dependence also varies more smoothly across GB for larger d (as many gauges will be the same in neighbouring sets) although the lower dependences in hilly areas are still apparent. For this value of T (equivalently p) estimates of $N_i(p, 90)$ and $N_i(p, 120)$ are approximately 0 for north-west GB, see Keef (2007).

The 95% confidence intervals for $N_i(p, 60)$, for $T = 55$ years, shown in Figures 7(a) and (c) were derived through block bootstrapping for the dependence estimates in Figure 7(b). The widths of the intervals are around 0.07, which suggests that the level of uncertainty is fairly small. The widths increase with increasing return period, and for $T = 548$ years they are around 0.12 (not shown). Unlike when extrapolating return levels the uncertainty in $N_i(p, 60)$ does not always increase with increasing T as $N_i(p, 60)$ is an expected proportion which tends to its lower bound of zero.

4.5 Spatial dependence in river flows

The results of the spatial dependence analysis of large river flows are presented in Figures 8 and 9. To define return periods it is necessary to specify a value for k in expression (1.1). The results in Sections 4.2 and 4.3 suggest that for many catchments $k = 2$ is realistic, but for some the return periods stated will be too low. Using $k = 2$ we estimated $N_{|i}(p, d)$ for p equal to 0.975, 0.995, 0.999, 0.9995, 0.9999, and 0.99995 which correspond to the $T = 0.2, 1, 5, 10, 50$ and 100 year return periods and for $d = 30$ and 60 km. The median numbers of neighbours considered were 3 for $d = 30$ km and 12 for $d = 60$ km.

Overall, the spatial dependences in river flows are stronger than for precipitation. This may be because river flow represents an areally averaged process, whereas the precipitation observations are made at a point and are therefore more vulnerable to small-scale variation,

for example from convection. Figure 8(c) shows the $N_i(p, 30)$ for p corresponding to a 50 year return period. A couple of locations have dependence estimates exceeding 0.7, with most of GB exceeding 0.3. For comparison consider for precipitation $N_i(p, 30)$ for $T = 55$ year, i.e. roughly the same return period, then Figure 6(c) showed that $N_i(p, 30)$ exceeded 0.4 at several locations and for the bulk of GB it exceeds 0.2.

The pattern of stronger dependence in the south-east GB compared with the north and west, which was seen for the precipitation analysis, is barely hinted at in Figures 8 and 9(b). A south-east to north-west pattern may be masked by the more pronounced local variation in the dependence in river flows. For example, an area in the Thames catchment in southern England stands out as having considerably lower dependence than the surrounding region. This shows up particularly clearly for $T = 0.2$ years and $d = 30$ km in Figure 8(a). In this area very permeable, and sometimes also larger than average, catchments are located close to more impermeable catchments (Figure 1). The slower flow response of the permeable catchments compared with the less permeable ones means that spatial dependence in river flow decreases when both types of catchments are found within $\Delta_i(d)$.

A couple of areas further north, in the Lake District and eastern Scotland, also have lower levels of dependence. This seems to be related to whether or not there are lakes and reservoirs in the catchments. The spatial dependence is particularly low for catchments that do not contain large water-bodies (Calder at Calder Hall, Brathay at Jeffy Knotts, Dighty Water at Balmossie Mill, Lyon at Comrie Bridge, Almond at Almondbank and Dean Water at Cookston), when the surrounding catchments within 30 km of them do. The effect of lakes and reservoirs is to attenuate the floods, so that the river levels rise and fall more slowly than they would otherwise do.

When there are several gauges on the same river within $\Delta_i(d)$, one would expect strong dependence in that area. However, because different geologies and other characteristics of the different sub-catchments affect the river flow formation, the effect of nested catchments is not especially obvious in the plots. A set of gauges along a river that do show strong dependence include the Dee at Woodend, the Dee at Park and the upstream tributary Girnock Burn at Littlemill, located in north-eastern Scotland. In general, conditioning gauges with large catchments show stronger dependence estimates than do conditioning gauges with small catchments,

presumably because the large catchments contain nested sub-catchments within $\Delta_i(d)$.

Similar to the results for the precipitation analysis, the spatial pattern of dependence in river flooding across Britain becomes smoother with increasing d (Figures 8(a) and 9(b)). It also tends to decrease with increasing T (Figure 8). Since precipitation is a major driver for river flow formation, this is probably largely related to the spatial structure of precipitation events, with more extreme falls embedded within areas of less extreme precipitation depths. However, it can be questioned whether flooding would ever occur, say, at a downstream gauge in a catchment without high flows occurring also at several of the upstream gauges. The model is fitted using records which are generally around 40 years in length, and extrapolation to higher return periods is based on the behaviour of these relatively frequent events (where dependence decreases with T). If there is indeed a change in the dependence structure for less frequent events longer records are needed for this to be picked up by the model.

Figures 9(a) and (c) show the lower and upper bounds of the 95% confidence intervals derived through block bootstrapping for the dependence estimates in Figure 9(b). The widths of the intervals are around 0.09, which is slightly larger than for the precipitation for the same T , but still rather small. The width increases with increasing T , and for $T = 100$ years they are around 0.12 wide (not shown). There is more variation in the confidence intervals for river flows than for precipitation. The widths of the intervals do not seem to change with the magnitude of the dependence estimates, but several widths exceed 0.2, particularly for the higher return periods. These larger confidence intervals seem to correspond to cases with a large amount of missing data for the conditioning gauge.

4.6 Factors affecting dependence in river flows

As shown in Section 4.5 the areas of the country with diverse catchments appear to exhibit lower spatial dependence of flooding. To investigate the effect of catchment characteristics on dependence (i.e. probability of widespread flooding) we carried out a correlation analysis to determine whether catchments with similar characteristics had higher dependence and vice versa. We looked for correlations between the absolute differences in catchment characteristics between two flow gauges and the pairwise dependence between these flow gauges. The data used in this study is a subset of the set used in the main study. In total, 36 conditioning gauges,

and gauges within 60 km of these conditioning gauges, were used for a pair-wise dependence analysis.

The catchment characteristics investigated include the catchment descriptors derived for the Flood Estimation Handbook (FEH) (Bayliss, 1999). In addition, the distance between each pair of gauges and the value of standard average annual rainfall at the conditioning gauge (CSAAR) were used. The latter was included as a proxy to see if the geographical variation in dependence in precipitation, which is strong (weak) in the south-east (north-west) with predominantly low (high) CSAAR values, is carried through to the dependence in flow.

Table 2 shows the Kendall's tau dependence measure between each of τ_{\max} , $P_{j|i}^{(\tau_{\max})}(p)$ and $Q_{j|i}^{(L)}(p)$ with the differences in each of the catchment descriptors. For assessing τ_{\max} , the differences in catchment descriptors were taken as the value for the non-conditioning gauge minus the value for the conditioning gauge, whereas for the other two analyses the absolute difference of the descriptors was used. The descriptors that have a noticeable effect on τ_{\max} include the catchment area, mean drainage path length, base flow index, standard percentage runoff, and flood attenuation by reservoirs and lakes. The same catchment descriptors are found to be of importance for both $P_{j|i}^{(\tau_{\max})}(p)$ and $Q_{j|i}^{(L)}(p)$, although the relationships tend to be stronger for the former. This seems reasonable, since the $P_{j|i}^{(\tau_{\max})}(p)$ allows the maximum dependence to be found among a larger number of different lags. All the descriptors except CSAAR have noticeable correlations with $P_{j|i}^{(\tau_{\max})}(p)$. However, two of the descriptors, the mean drainage path length and the flood attenuation by reservoirs and lakes, do not have a noticeable effect on $Q_{j|i}^{(L)}(p)$. The most influential differences for both dependence estimates are the differences in mean drainage path length, proportion of time when the soil is wet, standard average annual rainfall, base flow index, standard percentage runoff, and distance between the gauges.

5 Summary and concluding remarks

In this paper we have presented a complete study of within-event dependence of extreme river flows and precipitation over the whole of Great Britain. We have shown that it is possible to quantify the level of extremal dependence for a large number of sites and hence derive measures of spatial dependence over small and large spatial scales. The method that we have used in this paper is the first statistical method that is capable of producing reliable estimates of the

probabilities of multiple, simultaneous floods and spatially widespread extreme precipitation.

For both precipitation and river flows, events become more localised in space as the return periods of these events increases. For precipitation, spatial dependence is weaker in the mountainous north and west than in the rest of Great Britain, reflecting the effect of topographic variation. This effect is not obvious in the spatial dependence pattern for river flows, which shows a more variable pattern at smaller scales. The local variability is related to differences in catchment characteristics. In areas where characteristics are similar, spatial dependence is strong. However, where, for example, fast-responding impermeable catchments are located next to slowly-responding permeable catchments, the spatial dependence in river flows is low.

In this paper we ignore the seasonal features of the processes. However, Keef (2007) explored the seasonality of dependence in precipitation and found weaker spatial dependence in summer than winter. This may be attributed to the larger numbers of localised convective storms in summer than winter. In contrast, autumns and winters in Great Britain tend to be dominated by spatially widespread precipitation of frontal origin. Because of the seasonal changes in precipitation, it seems likely that spatial dependence in river flows may also be weaker in summer than in winter. Another feature of flow data that we have not examined is the effect of non-stationarity of very slowly responding catchments. The flows in these catchments are baseflow dominated and are dominated by changes in groundwater rather than recent precipitation. For these catchments the duration of high flow periods can last for weeks rather than days. Because the high flows of these catchments are so different from the high flows of most catchments further work is needed to ensure that the spatial aspects of flooding for these catchments are correctly included in any analysis.

Acknowledgements

CK thanks NERC for funding her studentship at Lancaster University and CEH Wallingford, reference number NER/S/A/2003/11343. We gratefully acknowledge the use of the daily mean river flow data, provided by the National River Flow Archive, and the daily precipitation totals, provided by the UK Met Office. We would also like to thank the referees for their helpful comments that have improved this paper.

References

- Bayliss, A. C., 1999. Flood Estimation Handbook Vol. 5: Catchment Descriptors. Institute of Hydrology, Wallingford.
- Coles, S. G., Heffernan, J. E., Tawn, J. A., 1999. Dependence measures for extreme value analyses. *Extremes* 2, 339–365.
- Dales, M. Y., Reed, D. W., 1989. Regional flood and storm hazard assessment. Tech. Rep. 102, Institute of Hydrology, Wallingford, UK.
- Environment Agency, 2009. Spatial coherence of flood risk, SC060088/SR1. Tech. rep., Environment Agency, Bristol.
- Heffernan, J. E., Tawn, J. A., 2004. A conditional approach for multivariate extreme values (with discussion). *J. R. Statist. Soc. B* 66, 1–34.
- Joe, H., 1997. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.
- Keef, C., 2007. Spatial Dependence of River Flooding and Extreme Rainfall. Ph.D. thesis, Lancaster University.
- Keef, C., Tawn, J. A., Svenson, C., 2009. Spatial risk assessment for extreme river flows. *Applied Statistics* 58 (5).
- Manley, G., 1970. Climate of the British Isles. In: Wallén, C. C. (Ed.), *Climates of Northern and Western Europe*, World Survey of Climatology Volume 5. Elsevier, London, UK, pp. 81–134.
- Marsh, T. J., Hannaford, J., 2008. UK Hydrometric Register. Hydrological Data UK Series., Centre for Ecology and Hydrology, Wallingford.
- URL <http://www.ceh.ac.uk/products/publications/UKHydrometricRegister.html>
- Mayes, J., Sutton, G., 1997. Eastern England. In: Wheeler, D., Mayes, J. (Eds.), *Regional Climates of the British Isles*. Routledge, London, UK, pp. 89–110.
- Nelsen, R. B., 1999. *An Introduction to Copulas*. Springer-Verlag, New York.
- Svensson, C., Jones, D., 2002. Dependence between extreme surge, river flow and precipitation in eastern Britain. *Int. J. Climatol.* 22 (10), 1149–1168.

Table 1: Catchment area, in km^2 , and BFI for the test river flow dataset. The first four are Scottish gauges, the last four are in the Thames catchment.

River	Location	Area	BFI
Tweed	Kingledores	139	0.45
Earn	Kinkell Bridge	591	0.50
Teith	Bridge of Teith	518	0.43
Tweed	Peebles	694	0.56
Thames	Kingston	9948	0.64
Thames	Eynsham	1616	0.68
Kennet	Theale	1033	0.87
Lambourn	Shaw	234	0.97

Table 2: Kendall's tau dependence measure between each of τ_{\max} , $P_{j|i}^{(\tau_{\max})}(p)$ and $Q_{j|i}^{(7)}(p)$ with the differences in catchment descriptors at the pair of gauges (i, j) . The catchment descriptors are: AREA - catchment area (km^2); DPL - mean drainage path length (km); DPS - mean drainage path slope (m/km); PROPWET - proportion of time when the soil is wet; SAAR - standard average annual rainfall (mm); BFI - base flow index; SPR - standard percentage runoff; FARL - flood attenuation by reservoirs and lakes; URBEXT - extent of urban and suburban land cover in 1990; CSAAR - value of SAAR at gauge i ; and DIST - the distance between each pair of gauges.

Covariate	τ_{\max}	$P_{j i}^{\tau_{\max}}(u)$	$Q_{j i}^{(7)}(p)$
AREA	0.345	-0.079	0.05
DPL DPS	0.344	-0.092	-0.009
PROPWET	0.047	-0.172	-0.116
SAAR	-0.009	-0.222	-0.184
BFI SPR	0.036	-0.239	-0.185
FARL	0.2	-0.183	-0.116
URBEXT	-0.189	-0.243	-0.202
CSAAR	-0.12	-0.09	-0.011
DIST	-0.02	-0.069	-0.073
	0.006	-0.013	-0.023
	-0.003	-0.201	-0.196

Figure 1: Locations of the selected river flow gauging stations. Characteristics of the stations are shown: baseflow index (left) and catchment area (km^2) (right).

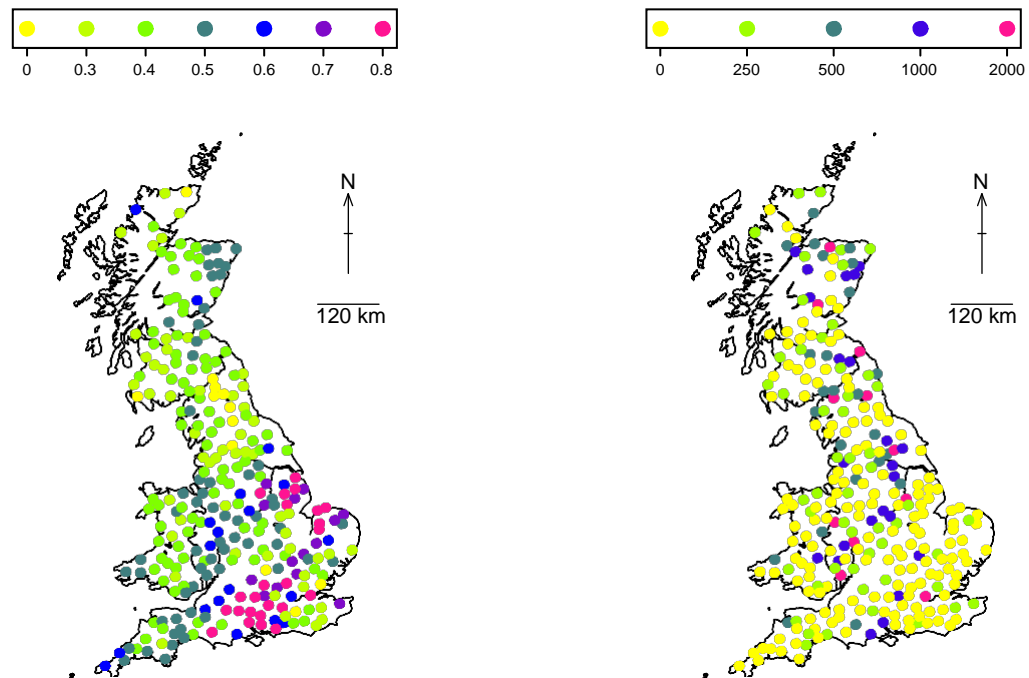


Figure 2: Locations of the selected raingauges. Characteristics of the gauge locations are shown: altitude (m) (left) and standard average annual rainfall (mm) (right).

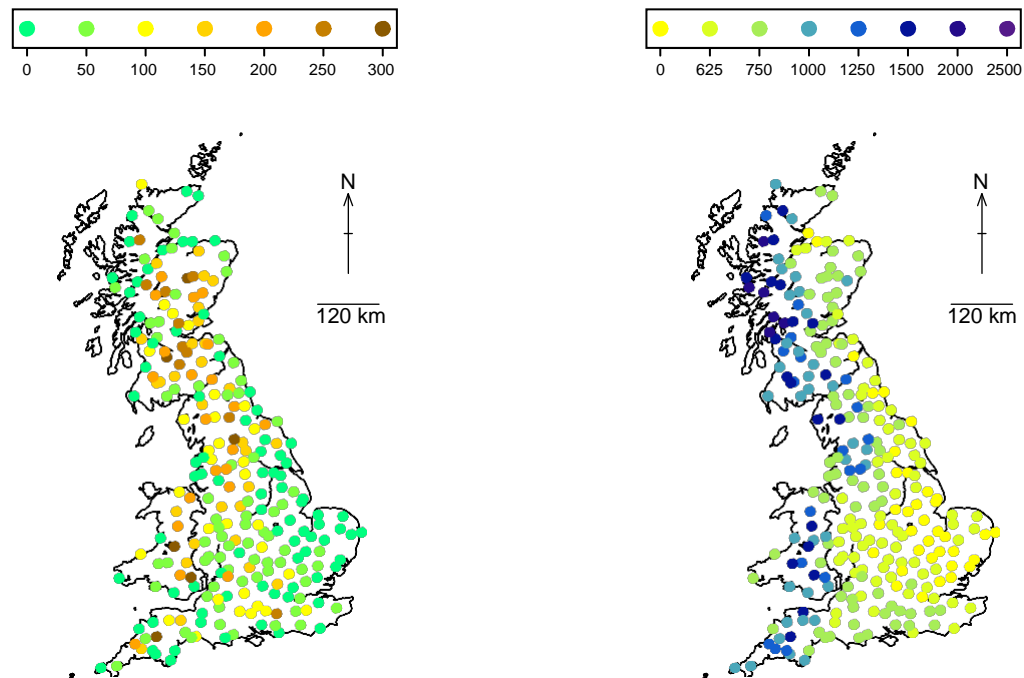


Figure 3: Top plots are data simulated for Heffernan and Tawn model, solid lines the conditional median and dashed lines the 0.025 and 0.975 conditional quantiles of $Y_j|Y_i, Y_i > 4.5$. Left plot $\alpha_{j|i} = 0.3, \beta_{j|i} = 0.7$, right plot $\alpha_{j|i} = 0.8, \beta_{j|i} = 0.1$. In both data sets $Z_{j|i}$ follows a Normal distribution with mean 0.5 and variance 0.25 and Y_i, Y_j follow a standard Gumbel distribution. Lower plots show $P_{j|i}^{(\tau)}(v) = \Pr(Y_j > v|Y_i > v)$ for the respective data sets, solid lines show true values, dashed lines show values estimated empirically from the simulated data, dotted lines show values estimated by fitting the model to the simulated data.

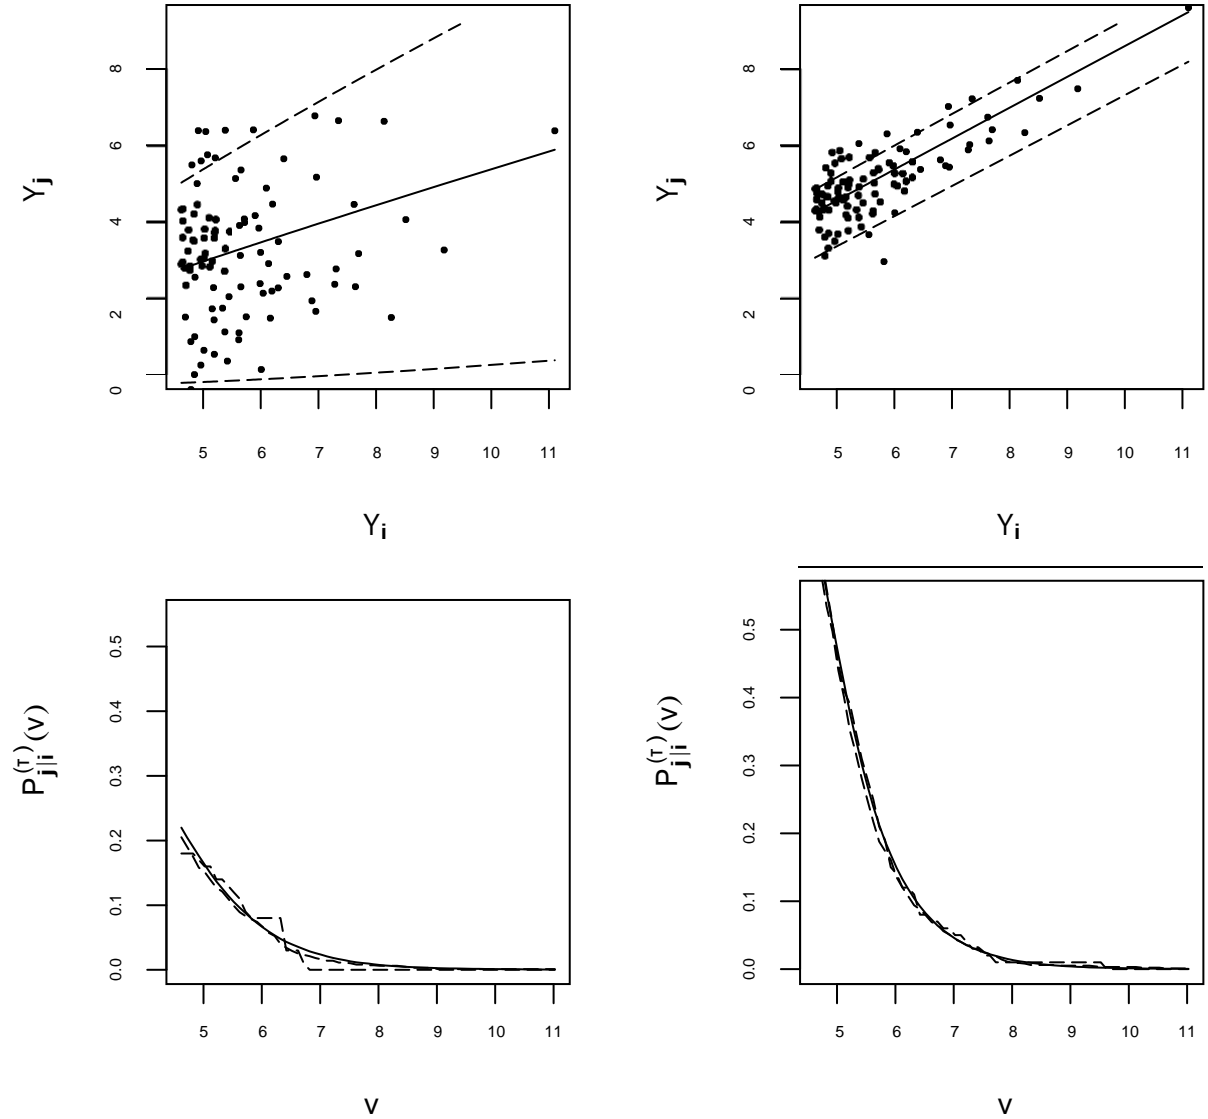



Figure 4: Temporal dependence in river flow at a) the Tweed at Peebles and b) the Thames at Eynsham, for lags from -50 to 50 days.

(a)



(b)





Figure 5: Cross-series temporal extremal dependence between a) the Tweed at Kingledores and the Earn at Kinkell Bridge, at lags of -50 to 50 days with respect to the Tweed at Kingledores, and b) the Thames at Kingston and the Lambourn at Shaw, at lags of -50 to 50 days with respect to the Thames at Kingston.

(a)



(b)




Figure 6: Estimated spatial dependence in precipitation, $N_i(p, 30)$ corresponding to a range of T year return periods: (a) $T = 0.1$, (b) $T = 5$, (c) $T = 55$ and (d) $T = 548$, in all $L = 0$.

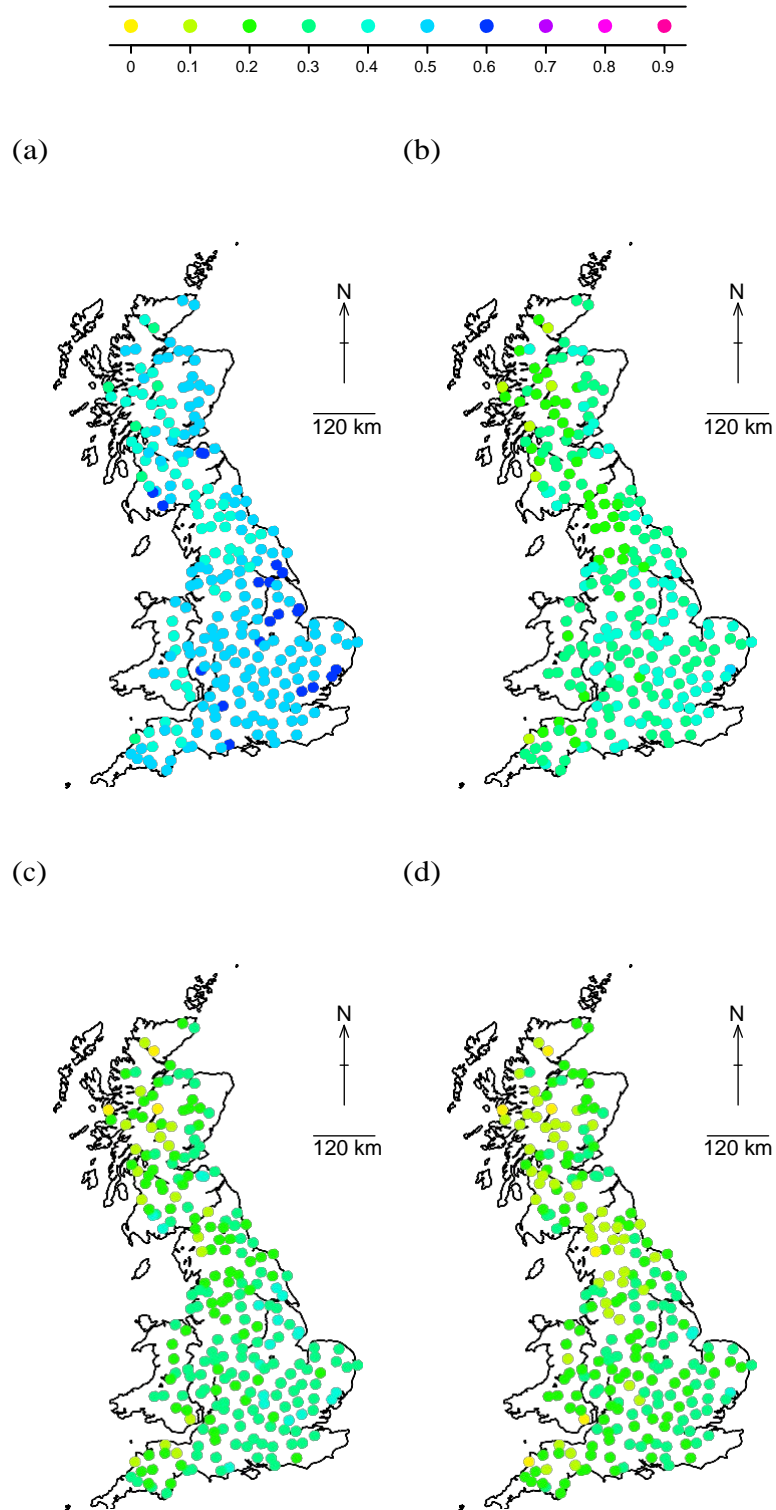


Figure 7: Estimated spatial dependence in precipitation, $N_i(p, 60)$ corresponding to $T = 55$ years: (a) lower bound of 95% confidence interval, (b) estimate, and (c) upper bound of 95% confidence interval, in all $L = 0$.

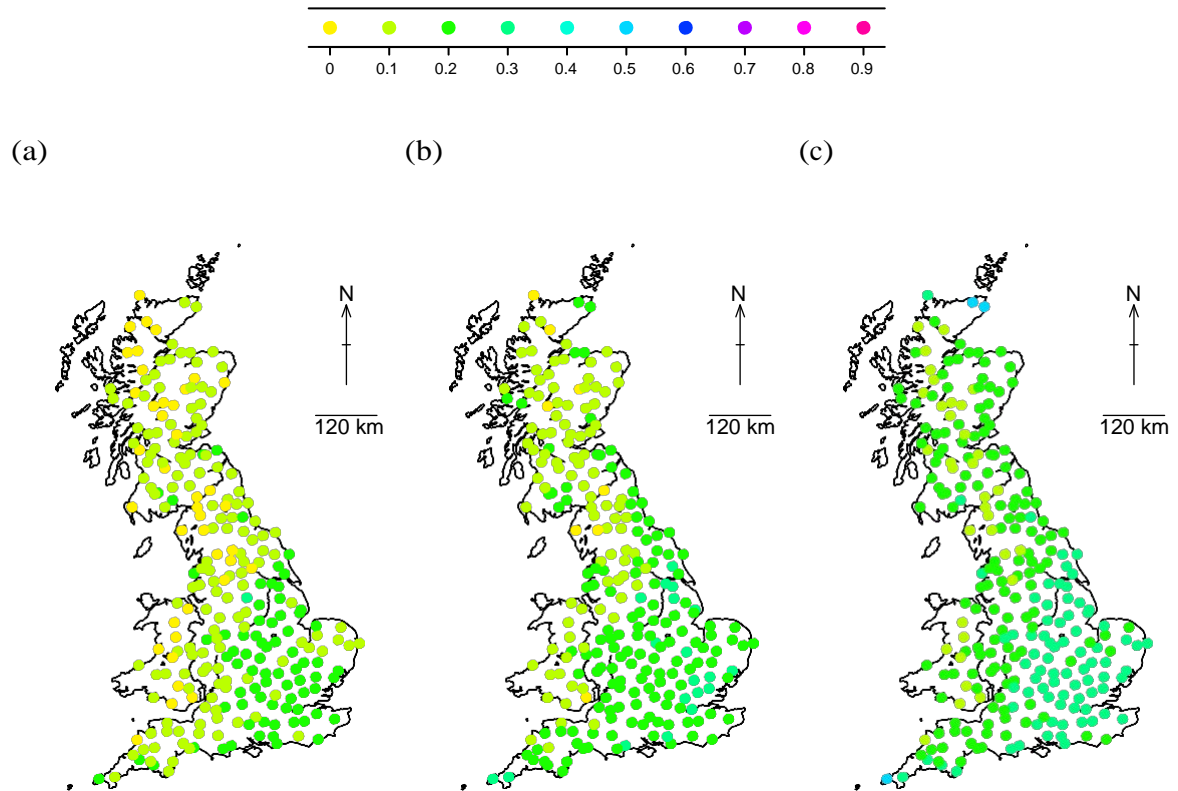


Figure 8: Spatial dependence in river flows, $N_i(p, 30)$: corresponding to a range of T year return periods: (a) $T = 0.2$, (b) $T = 5$, (c) $T = 50$ and (d) $T = 100$, in all $L = 3$.

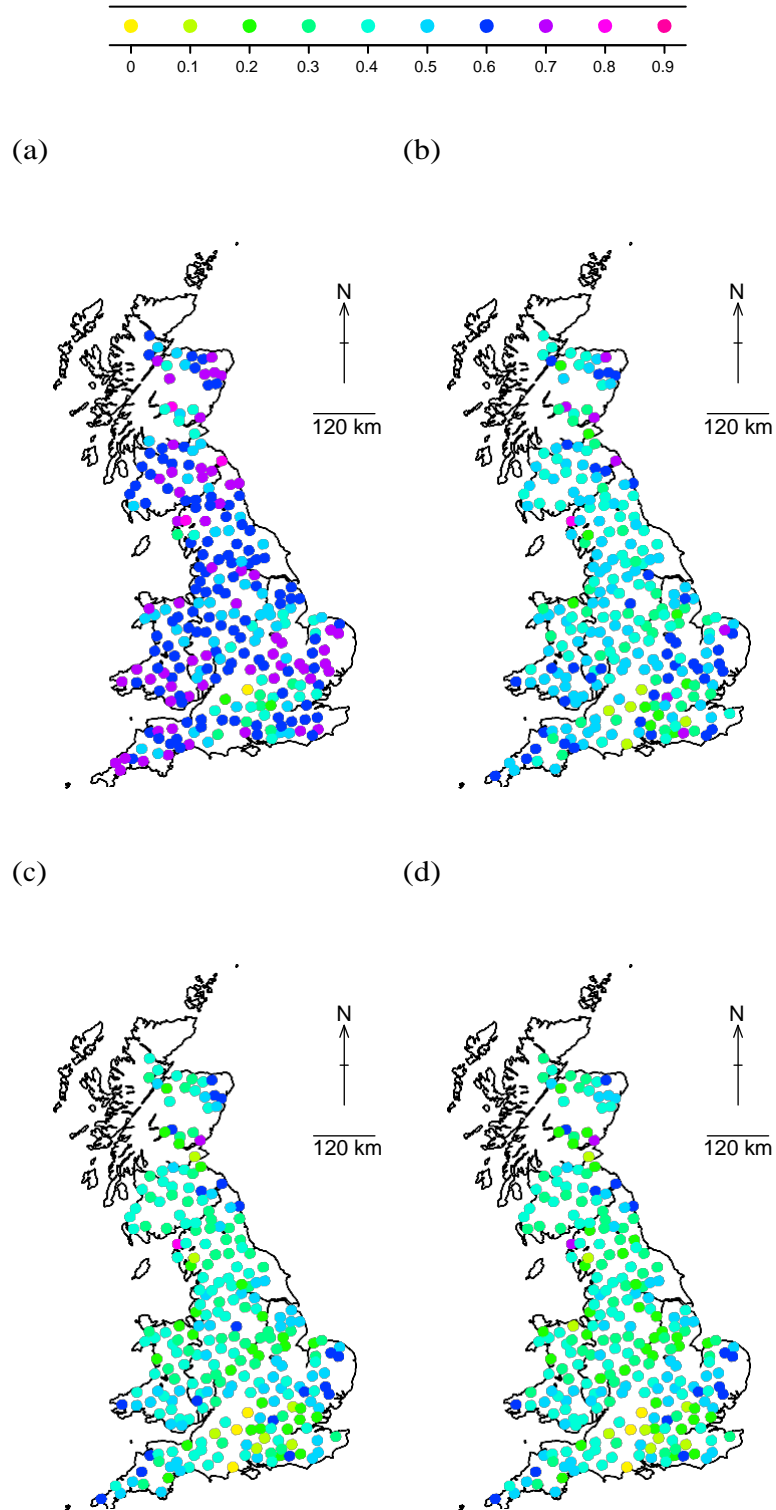


Figure 9: Estimated spatial dependence in river flows, $N_i(p, 60)$ corresponding to $T = 50$ years: (a) lower bound of 95% confidence interval, (b) estimate, and (c) upper bound of 95% confidence interval, in all $L = 3$.

